Essays on causal machine learning for decision-making

Causal learning integrates the fields of causal inference and machine learning to assess the impact of decisions on variables of interest to organizations. It aims to translate data into actionable insights while taking into account that certain actions affect each individual differently. As such, it serves as a basis for informed decisions. Unlike classification tasks that predict the future value of an outcome, causal models estimate the magnitude of the change in the value of an outcome due to the application of a treatment. This allows decision-makers to develop customized campaigns in which only the most favorable action is applied to individuals.

The task of anticipating the individual-level effect of a decision on an outcome is known as Individual Treatment Effect (ITE) estimation. A treatment is a variable under the control of a decision-maker whose manipulation is expected to influence the behavior of individuals. The prediction of ITEs have been facilitated by advances in the field of machine learning. This evolution considerably extended the field of treatment effects estimation, which is common in econometric studies. Particularly, it broadened the scope of causal inference beyond the evaluation of the effect of treatments at the population-level, towards a more active role in the design of personalized treatments.

ITE estimation is not limited to single treatment applications. Practical cases involve predicting the ITEs of various treatment alternatives. Hence, we survey the literature on multitreatment causal learning and present a classification of multitreatment causal models. Two novel approaches, the naive uplift approach and the multitreatment modified outcome approach, are proposed for multitreatment ITE estimation. We present the results of a benchmarking study and show that none of the evaluated techniques consistently outperform other techniques. However, the newly proposed techniques are found to offer similar performances compared to state-of-the-art approaches.

We argue that causal models should not only be evaluated on the basis of the expected gain in terms of treatment effects. The decision of offering and applying a treatment entails benefits and costs. Hence, the profit generated by a causal model is an essential metric to assess model performance. We integrate the cost-sensitive classification framework with causal learning from which a business-oriented performance metric is derived: causal profit. The proposed causal cost-sensitive framework encompasses previous causal cost-sensitive approaches in the literature on causal customer retention and causal customer response.

We derive the causal cost-sensitive decision boundary and the expected causal profit ranker from the causal cost-sensitive framework. The causal cost-sensitive decision boundary acts as a decision rule from which treatment is prescribed. It is based on the estimated ITE, the positive outcome probability of treatment, and cost and benefit parameters of the problem setting. In addition, the causal profit ranker maximizes the expected profit across all possible targeting thresholds and prioritizes instances for whom treatment should be prescribed.

This dissertation also presents an application of causal learning in educational data mining. We demonstrate that causal models can improve the effectiveness of retention efforts in institutions of higher education. Various causal models are trained to estimate the ITEs of a tutoring program. The results indicate that a causal model outperforms conventional predictive modeling approaches in terms of the expected treatment effect.

The dissertation concludes by discussing the main findings and outlining future research opportunities. Causal learning from observational data heavily relies on assumptions, which in turn have influenced our findings. Therefore, an important step towards popularizing the use of causal models is to develop tools that provide transparency on the structure of the causal relationships and facilitate the evaluation of assumptions, e.g., causal graphs. Moreover, causal relationships are dynamic, and individuals are becoming increasingly interconnected. These motivate the assessment of the effect of time and network interference in the estimation of ITEs.

In summary, the doctoral dissertation contributes to the field of causal learning. First, it proposes new tools to estimate ITEs in multiple treatment cases. Second, it extends the cost-sensitive classification framework to evaluate the performance of causal models. Third, it develops the causal cost-sensitive decision rule and causal profit ranker for causal cost-sensitive decision-making. Lastly, it also contributes to the field of educational data mining. The case study on student dropout demonstrates the potential of causal models in designing more efficient tutoring programs.